

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project 0704-0188, Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED FINAL - 30 SEP 92 TO 30 SEP 95
4. TITLE AND SUBTITLE STRATEGIES FOR THE CONTROL OF REASONING IN DYNAMIC ENVIRONMENT		5. FUNDING NUMBERS F49620-92-J-0422 2304/GS 61102F
6. AUTHOR(S) PROFESSOR MARTHA E. POLLACK		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) DEPARTMENT OF COMPUTER SCIENCE UNIVERSITY OF PITTSBURGH PITTSUBRH, PA 15260		
8. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NM 110 DUNCAN AVE, SUITE B115 BOLLING AFB, DC 20332-0001		AFOSR-TR-96 OICD F49620-92-J-0422

12a. DISTRIBUTION AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED	12b. DISTRIBUTION CODE
12c. ABSTRACT (MAXIMUM 250 WORDS) SEE ABSTRACT ON REPORT.	

19960320 065

UNCLASSIFIED

UNCLASSIFIED

UNCLASSIFIED

SAR

RIA 25/10/97

STRATEGIES FOR THE CONTROL OF REASONING IN DYNAMIC ENVIRONMENT

Final Report

November 22, 1995

Prof. Martha E. Pollack
Dept. of Computer Science
Univ. of Pittsburgh
Pittsburgh, PA 15260

Prepared for:
Air Force Office of Scientific Research
AFOSR, Bolling Air Force Base
Washington, D.C. 20332-6448
Attn: Dr. Abraham Waksman

1 Introduction

This document summarizes the research conducted on the AFOSR-supported project "Strategies for the Control of Reasoning in Dynamic Environments," Contract F49620-92-J-0422, during the period between Oct. 1, 1993 and Sept. 30, 1995. The goal of the project has been to investigate strategies that enable intelligent, real-time problem solving systems (IRTPS) to control their reasoning resources in dynamic, unpredictable environments. Our efforts have been focused in four key areas:

- analyzing the requirements that must be met by simulation systems to support the analysis of reasoning control strategies, and developing simulators that meet those requirements;
- developing and analyzing strategies for controlling deliberation (reasoning about alternatives) in dynamic environments;
- developing and analyzing strategies for controlling planning (means-end reasoning) in dynamic environments;
- developing and analyzing strategies for controlling interactions in multi-agent dynamic environments.

This report is organized around these four topics: after a brief overview of the research problem, there is a section devoted to our research results on each topic. This is followed by list of project-sponsored papers and a summary of students supported by this project. Copies of project-supported papers are included as an appendix.

2 Overview

As noted above, the overall goal of the project has been to investigate strategies to enable intelligent, real-time problem solving systems (IRTPS) to control their reasoning resources in dynamic, unpredictable environments. Typically, the designers of such systems painstakingly hand-encode procedures to manage reasoning in response to those domain-specific conditions that can be anticipated. System development is thus extremely labor intensive, and it is difficult

to guarantee the robustness or reliability of the resulting systems. These problems would be alleviated by a better understanding of the relationships among alternative system designs and reasoning strategies (e.g., a tendency not to reconsider once a goal is adopted), the constraints of the system's physical capabilities (e.g., the accuracy of its sensors), and the characteristics of the system's intended environment (e.g., the rate at which it is changing).

Reliable, intelligent, real-time performance in dynamic environments is difficult to achieve, because all systems have inherent computational limits: they cannot perform arbitrarily large computations in finite time. While a system in a dynamic environment is reasoning about what actions to perform, the environment may change. In fact, it may even change in ways that undermine the assumptions underlying the ongoing reasoning. A system may begin a deliberation process with a particular set of available options, but in a dynamic environment, new options may arise and formerly existing options disappear during the course of the deliberation. Moreover, the utilities associated with each option are subject to change during the course of the deliberation. A system that blindly pushes forward with its original deliberation process, without regard to the amount of time it is taking or the changes meanwhile going on, is not likely to make rational decisions.

There is thus a need to develop strategies that will allow an agent to make efficient decisions about what reasoning problems to focus on when, and for how long. These control strategies must enable the agent to trade potential decision quality for decision timeliness. A related trade-off involves the cost of communication, which must be weighed against the potential benefit of additional information gained as a result of the communication. In some situations, it is advantageous for an agent to delay action in order to communicate with other agents, while in other circumstances communication will not be worth the cost.

In the current project we have been developing and analyzing strategies for controlling reasoning and communication, in single-agent and in multi-agent environments. We have studied two types of reasoning-control strategies: those that enable the agent to manage its deliberation (i.e., its reasoning about what options to pursue), and those that enable it to manage its planning (i.e., its reasoning about how to achieve the selected options). To investigate the relative advantages and disadvantages of particular strategies, we have employed an experimental methodology, conducting experiments using abstract simulation systems.

3 Experimental Platforms

Our research on this project has been done primarily using two simulation systems: the Tileworld system and the DIPART system. The Tileworld is an abstract dynamic environment with an embedded agent, an early prototype of which was developed prior to the start of this project; during the project itself, we made a number of enhancements to the system to enable it to support our experiments. The DIPART system is somewhat more closely connected to a real application, namely a noncombatant evacuation operation.

Our accomplishments on this topic include the following:

- careful analysis of the requirements of testbed systems for studying the control of reasoning; see references 9, 16, and 19.
- design, implementation, and testing of extensions to the Tileworld testbed system, to enable it to support a wider range of experiments; see references 22 and 13.
- adaptation of the Tileworld to support experiments in multi-agent environments; see reference 4.
- design, implementation, and testing of the DIPART testbed system; see references 23 and 27.¹

4 Single-Agent Deliberation-Control Strategies

The single-agent deliberation-control strategies that we studied can be classified into two groups: those that derive from our earlier work on resource-bounded reasoning and make central use of the concept that intentions can be used to “filter” reasoning options, and those that make use of techniques from the real-time operating-systems literature, particularly the notion of value-density based scheduling.

¹The development of DIPART has been largely supported by the Rome Laboratory (RL) of the Air Force Material Command and the Defense Advanced Research Projects Agency “Distributed, Interactive Development and Monitoring of Transportation Plans in Dynamic Environments” (Contract F30602-93-C-003).

Accomplishments on this topic include:

- development of a formal model of intention, which can serve as the theoretical foundation for filtering and related intention-based strategies; see references 14 and 21.
- detailed analysis of the relative effectiveness of the filtering strategy in different environments, including identification of conditions under which the filtering control strategy must be augmented with an override mechanism to prevent a decay in performance; see references 20 and 22.
- adaptation of value-density based scheduling algorithms to the control of deliberation in dynamic environments; see references 24 and 25.

5 Single-Agent Planning-Control Strategies

Meta-level control of reasoning is also important *during* the process of plan formation: that is, once an agent has decided to focus on a particular planning problem, major questions remain about how it should structure its reasoning during the search for a plan. Our accomplishments on this topic include:

- development of the least-cost flaw repair (LCFR) strategy for controlling planning; experimental validation that LCFR reduces search-space size; and development of an algorithm that, by approximating LCFR, reduces search time as well as space; see references 10 and 11.
- development and experimental validation of an algorithm for effectively controlling planning in situations in which goals can be partially satisfied; see reference 15.
- development and experimental validation of a new approach to planning, which supplements traditional refinement-search with constraint-satisfaction techniques, thereby enabling finer-grained control of the planning process, with concomitant efficiency gains; see references 10 and 12.
- development and implementation of a planning algorithm that combines causal and decompositional reasoning; see reference 26.

- preliminary work on the question of controlling conditional planning, i.e., determining which contingencies should be handled at plan time and which should be deferred until execution time; see reference 17.

6 Multi-Agent Coordination Strategies

Multi-agent environments are inherently dynamic, and hence strategies for controlling reasoning are essential in them. However, these environments pose additional challenges, because of the need for agents not only to control their own reasoning but also to coordinate with one another. Our accomplishments on this topic include:

- generalization of the filtering strategy to the multi-agent case, and analysis of the conditions under which it is effective; see reference 4.
- development of multi-agent planning algorithms for various models of interaction, based on the aggregation of individually generated plans; see references 1, 2, 3, 6, and 7.
- specification of fast techniques to derive consensus and coordination in multi-agent environments; see references 5 and 8.

7 Project-Supported Papers

This list provides all the project-supported papers, which are referred to in the three previous sections. Complete bibliographic references to related work can be found in the papers themselves, all of which are included in an appendix to this report.

1. E. Ephrati, M. Perry, and J. S. Rosenschein. Plan execution motivation in multi-agent systems. *Proceedings of the Second International Conference on Artificial Intelligence Planning Systems (AIPS)*, June, 1994.
2. E. Ephrati, M. E. Pollack, and J. S. Rosenschein. Exploitation of decision theory techniques in multi-agent planning. *1994 AAAI Spring Symposium on Decision Theoretic Planning*, AAAI Press, March, 1994.

3. E. Ephrati, M. E. Pollack, and J. S. Rosenschein. A tractable heuristic that maximizes global utility through local plan combination. *Proceedings of the 1st International Conference on Multi-Agent Systems (IC-MAS)*, June, 1995.
4. E. Ephrati, M. E. Pollack, and S. Ur. Deriving multi-agent coordination through filtering strategies. *Proceedings of the 14th Joint International Conference on Artificial Intelligence (IJCAI)*, August, 1995.
5. E. Ephrati and J. S. Rosenschein. Deriving consensus in multi-agent systems. To appear in *Artificial Intelligence*.
6. E. Ephrati and J. S. Rosenschein. A heuristic technique for multiagent planning. To appear in *Annals of Mathematics and Artificial Intelligence*.
7. E. Ephrati and J. S. Rosenschein. Divide and conquer in multi-agent systems. *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, August 1994.
8. E. Ephrati, G. Zlotkin, and J. S. Rosenschein. Meet your destiny: a non-manipulable meeting schedule. *Proceedings of the ACM 1994 Conference on Computer Supported Cooperative Work*, October, 1994.
9. S. Hanks, M. E. Pollack, and P. R. Cohen. Benchmarks, testbeds, controlled experimentation, and the design of agent architectures. *AI Magazine*, 14(4):17–42, 1993.
10. D. Joslin. *Passive and Active Commitment in Plan Generation* Ph.D. Dissertation, Department of Computer Science, Univ. of Pittsburgh, December 1995.
11. D. Joslin and M. E. Pollack. Least-cost flaw repair: A plan refinement strategy for partial-order planning. *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, August 1994.
12. D. Joslin and M. E. Pollack. Active and passive decision postponement in plan generation. *Proceedings of the Third European Workshop on Planning (EWSP)*, September, 1995.

13. D. Joslin, A. Nunes and M. E. Pollack. *TileWorld Users' Manual*. University of Pittsburgh Tech. Report TR 93-12, August, 1993.
14. K. Konolige and M. E. Pollack, A Representationalist Theory of Intention. *Proceedings of the 13th International Joint Conference on Artificial Intelligence (IJCAI)*, August 1993.
15. M. Milshtein. *A Cost-Directed Heuristic Planner*. M.S. Project Report, Dept. of Computer Science, University of Pittsburgh, October, 1995.
16. A. Nunes. *Towards a Machine IQ: Some perspectives on evaluation metrics for embedded agents with an application to agents designed for the Pacifica evaluation scenario*. M.S. Project Report, Dept. of Computer Science, University of Pittsburgh, May, 1994; Available as Tech. Report 94-28.
17. N. Onder. *What contingencies to plan for*. Ph.D. dissertation proposal, Dept. of Computer Science, University of Pittsburgh, March, 1995; Available as Tech. Report 95-25.
18. M. E. Pollack. Review of "Artificial Intelligence: A Modern Approach." *AI Magazine* 16(3): 73-74, 1995.
19. M. E. Pollack. Evaluating plans, planners, and planning systems. *SIGART Bulletin* 6(1), 1995.
20. M. E. Pollack. Filtering as a reasoning-control strategy: An experimental assessment. In *Proceedings of SOAR '93 (Space Operations and Applications, and Research Symposium)*, 1993.
21. M. E. Pollack. Review of "Agency in Action: The Practical Rational Agency Machine." *Computational Linguistics*, 19(3):565, 1993.
22. M. E. Pollack, D. Joslin, A. Nunes, S. Ur, and E. Ephrati. *Experimental Investigation of an Agent Commitment Strategy*. Dept. of Computer Science, Univ. of Pittsburgh, Tech. Report 94-31; under revision for publication.
23. M. E. Pollack, T. Znati, E. Ephrati, D. Joslin, S. Lauzac, A. Nunes, N. Onder, Y. Ronen, S. Ur, "The DIPART Project: A Status Report." *Proceedings of the ARPI Planning Initiative Workshop*, Tucson, AZ, 1994.

24. Y. Ronen. *Meta-Level Deliberation as Scheduling: The Use of Operating Systems and Operations Research Techniques in Meta-Level Control*. M.S. thesis, Dept. of Computer Science, Univ. of Pittsburgh, Aug., 1995.
25. Y. Ronen and M. E. Pollack. Value-density algorithms for deliberation scheduling. *Proceedings of the IJCAI Workshop on Anytime Algorithms and Deliberation Scheduling*, August, 1995.
26. R. M. Young, , M. E. Pollack, and J. D. Moore. Decomposition and causality in partial-order planning. *Proceedings of the Second International Conference on Artificial Intelligence Planning Systems (AIPS)*, June, 1994.
27. T. Znati and M. E. Pollack. DIPART, an interactive simulation platform for plan development and monitoring in dynamic environments. *Proceedings of the 27th Annual Simulation Society*. IEEE Computer Society Press, April, 1994, pages 236-244.

8 Student Support

The following student theses were partially supported by this contract:

1. **David Joslin**, Ph.D., December, 1995. Dissertation: *Passive and Active Commitment in Plan Generation*. David is currently employed at the Computational Intelligence Research Laboratory (CIRL), University of Oregon, Eugene, OR.
2. **Arthur Nunes**, M.S., May, 1994. M.S. Project: *Towards a Machine IQ: Some perspectives on evaluation metrics for embedded agents with an application to agents designed for the Pacifica evaluation scenario*. Arthur is currently a Ph.D. student in the Applied Mathematics department at the University of Texas at Arlington.
3. **Yagil Ronen**, M.S., August, 1995. M.S. Thesis: *Meta-Level Deliberation as Scheduling: The Use of Operating Systems and Operations Research Techniques in Meta-Level Control*. Yagil is currently employed at the Inference Corporation, New York.

4. **Marina Milshtein**, M.S. in progress. M.S. Project: *A Cost Directed Heuristic Planner*. Marina recently completed her M.S. project, and is now completing her coursework for her M.S. degree, which she should receive in May, 1996.